View metadata, citation and similar papers at core.ac.uk

brought to you by CORE

provided by DigitalCommons@University of Nebraska

University of Nebraska - Lincoln
DigitalCommons@University of Nebraska - Lincoln

CSE Journal Articles | Computer Science and Engineering, Department of

2008

Investigating adaptive, confidence-based strategic negotiations in complex multiagent environments

Leen-Kiat Soh
University of Nebraska - Lincoln, lsoh2@unl.edu

Xin Li
University of Nebraska-Lincoln, xinli@cse.unl.edu

Follow this and additional works at: http://digitalcommons.unl.edu/csearticles

Part of the Computer Sciences Commons

Soh, Leen-Kiat and Li, Xin, "Investigating adaptive, confidence-based strategic negotiations in complex multiagent environments" (2008). *CSE Journal Articles*. 100.
http://digitalcommons.unl.edu/csearticles/100

This Article is brought to you for free and open access by the Computer Science and Engineering, Department of at DigitalCommons@University of Nebraska - Lincoln. It has been accepted for inclusion in CSE Journal Articles by an authorized administrator of DigitalCommons@University of Nebraska - Lincoln.

Published in *Web Intelligence and Agent Systems: An International Journal* 6 (2008), pp. 313–326;
doi: 10.3233/WIA-2008-0143 Copyright © 2008 IOS Press. Used by permission.

Investigating adaptive, confidence-based strategic negotiations in complex multiagent environments

Leen-Kiat Soh and Xin Li

Department of Computer Science and Engineering, University of Nebraska–Lincoln,
256 Avery Hall, Lincoln, NE 68588-0115, USA

Corresponding author — L.-K. Soh, email lksoh@cse.unl.edu

Abstract
We propose an adaptive 1-to-many negotiation strategy for multiagent coalition formation in complex environments that are dynamic, uncertain, and real-time. Our strategy deals with how to assign multiple issues to a set of concurrent negotiations based on an initiating agent's confidence in its profiling of its peer agents. When an agent is confident, it uses a packaged approach—conducting multiple *multi-issue* negotiations—with its peers. Otherwise, it uses a pipelined approach—conducting multiple *single-issue* negotiations—with its peers. The initiating agent is also capable of using both approaches in a hybrid, dealing with a mixed group of responding peers. An agent's confidence in its profile or view of another agent is crucial, and that depends on the environment in which the agents operate. To evaluate the proposed strategy, we use a coalition formation framework in a complex environment. Results show that the proposed strategy outperforms the *purely pipelined* strategy and the *purely packaged* strategy in both efficiency and effectiveness.

Keywords: strategic negotiations, adaptive, confidence, negotiation management

1. Introduction

Negotiation is a form of interaction among autonomous agents in which a group of agents with a desire to cooperate but with potentially conflicting interests seek to reach an agreement on a set of issues [5,12,20]. A negotiation *issue* can be any good, service, scarce resource, or unique capability that one agent can provide to another agent for a reward. A negotiation may address multiple issues or only one issue at a time. In this paper we focus on *multi-issue negotiations* on task and resource allocation in cooperative problem solving.

An agent that encounters a problem may initiate 1-to-many negotiations and concurrently negotiate with multiple agents over resources or services (capabilities). During negotiations, both the agent that initiates negotiations (i.e., the initiating agent) and each peer agent that responds to a negotiation request (i.e., the responding agent) may need to use short-term tactics guided by long-term strategies to reach an agreement acceptable to both sides. This need is particularly prominent in resource-constrained, complex environments: any misuse of negotiation resources for solving a present task could potentially prevent the agent from handling subsequent tasks. Conducting effective and efficient negotiations in complex environments is difficult. Complex environments that we consider in this paper are dynamic, uncertain, and real-time. The initial states that trigger the agents' negotiation process in the first place may *dynamically* change while the negotiation process is still on-going. The negotiation outcome is *uncertain* such that a negotiation carried out in the same manner does not always yield the same results at different times. Also, each negotiation is *time-constrained* such that an agreement reached past a deadline specified for a task is not acceptable. If an agent simply reacts when it negotiates in such an environment, it will not be able to plan to conserve its

resources for later tasks. If an agent only looks at the big picture, it might not be able to meet the individual needs and behaviors of the different agents encountered during negotiations. The problem is compounded when there are multiple issues at hand and an initiating agent needs to conduct multiple, concurrent negotiations with more than one responding agent. Should the initiating agent negotiate with as many responding agents as possible as long as it has the resources to do so, thus possibly starving itself of those resources if a new task appears? Should the initiating agent identifies the best responding agent and tries to persuade that agent to take care of all issues, with the risk that if the negotiation falls through, the initiating agent would have to start the process with another responding agent, thus losing valuable time? What should the strategy be?

We propose an adaptive, confidence-based negotiation strategy, first described in [17]. Briefly, after identifying a set of capable peers for a particular task at hand, the initiating agent then decides how to conduct its multiple, concurrent negotiations. If the initiating agent is confident in a peer's consistency, then it uses a packaged approach; otherwise, a pipelined approach. In a *pipelined* approach, the two negotiating agents negotiate on one issue at a time. In a *packaged* approach, the two agents negotiate on multiple issues during one negotiation. Using the pipelined approach, the initiating agent can switch remaining, yet-to-be-negotiated issues to other peers that have completed their negotiations with the initiating agent. This allows the initiating agent to be more efficient and effective. With the packaged approach, the initiating agent has the advantage of likely reaching agreements sooner on all the pertinent issues of its task at hand, and using fewer resources for negotiation. Therefore, the proposed strategy chooses to use the pipelined approach if the initiating agent is not confident in its profile of its peers, and the packaged approach if it is. If the agent has a mixed set of confidence in its peers, then it chooses a hybrid set of pipelined and packaged approaches.

To evaluate the adaptive, confidence-based negotiation strategy, we incorporate the strategy into a Multi-Phase Coalition Formation (MPCF) framework [15,16]. Our experiments compare the performance of the adaptive, confidence-based negotiation strategy with the *purely pipelined* strategy and the *purely packaged* strategy and show that the confidence-based one outperforms the others. We also investigate how our strategy affects the process (efficiency) and the outcome (effectiveness) of negotiations under different environmental conditions.

Note that the focus of this paper is on *strategic negotiations* instead of detailed negotiation tactics. Specifically, we focus on the initiating agent's negotiation strategy on how to conduct concurrent 1-to-many negotiations with multiple issues. Here, a negotiation strategy does not deal with how each negotiation between agents should proceed at each interaction step; instead, it manages how the concurrent negotiations should be conducted in terms of scheduling and assigning different issues to different peer agents.

Note also that this paper is a significant extension of [17], with more detailed discussions on the design of the proposed strategy and additional experiments and insights. The framework of the learning-based Multi-Phase Coalition Formation was first proposed in [15] and then further described in [16]. In this paper, we briefly present the framework in Section 3 to provide a context for the proposed strategy.

2. Related work

The research on negotiations can be classified into two categories: *tactical* negotiations and *strategic* negotiations. Tactical negotiations research focuses on how to conduct a negotiation with the step-by-step negotiation tactics. The typical approaches include (1) the game-theoretic and auction-based mechanisms (e.g., [13]) that focus on analyzing the interaction as a game between identical participants, and seeking its equilibrium, (2) the heuristic-based bargaining mechanisms (e.g., [3]) that are used in cases where it is not possible to reach the optimal outcome due to resource limitations, dynamic environment or incomplete information, and are mainly based on empirical testing and evaluation, and (3) the argumentation-based approaches (e.g., [8]) that allow agents to exchange meta-information about proposals and counterproposals. Since this paper focuses on strategic negotiation, the following discussions are primarily on strategic and hybrid negotiations.

2.1. Strategic negotiation

In general, strategic negotiation addresses the management of negotiations such as conflict resolution, resource allocation, preference concerns, concurrent negotiation coordination, and so on. In some applications where agents are not completely autonomous, negotia-

tions between agents have to be coordinated by a central arbitrator to resolve the conflict (e.g., [19]). In some applications where there are resource constraints, strategic negotiations are employed to allocate resources more efficiently (e.g., [6,7]). In some applications where agents have special preferences like optional negotiation order, strategic negotiations are employed to address the preferences (e.g., [4]). In our applications we employ strategic negotiation to manage multiple concurrent negotiations.

Sycara [19] presents a model of strategic negotiation that combines case-based reasoning and optimization of the multi-attribute utilities of agents, to resolve the adversarial conflicts in the domain of labor relations. An interesting component here is the use of a centralized arbitrator to handle the disagreements between agents. This centralized arbitrator essentially acts as the strategist as it has the knowledge of all disagreements and the goings-on in the environment, and thus has the capacity to match-make or even coordinate negotiations to resolve disagreements. However, it is a bottleneck, and its view of the environment may be too broad to effectively address an agent's needs. Our proposed solution does not use such a centralized arbitrator. The agents make decisions autonomously in a peer-to-peer manner. Moreover, our strategy is aimed to compromise the advantages of different negotiation approaches to improve both effectiveness and efficiency in negotiation, rather than simply seek the optimal negotiation outcome.

Kraus et al. [6,7] study strategic negotiation to improve the negotiation efficiency, specifically the negotiation time, similar to our work. They focus on the passage of time and the preferences of the players for different agreements as well as for opting out of the negotiations. They assume agents have complete information about others. Different from their work, our confidence-based negotiation strategy focuses on the management of concurrent negotiations. Agents try to reduce the overall negotiation time by choosing appropriate negotiation approaches, instead of by reducing the step-by-step negotiation time. Agents are aimed to improve both negotiation effectiveness and efficiency at the same time. Moreover, our agents only have incomplete information about others. They take the strategic negotiations based on the confidence in negotiation partners' consistent behaviors instead of the preference, to adapt to the complex environment.

Most similar to our work, Fatima et al. [4] focus on bilateral multi-issue strategic negotiation. The outcome of negotiation depends on two key factors: the *agenda* (i.e., the set of issues under negotiation) and the *negotiation procedure* (i.e., whether the issues are discussed together or separately). The authors determine the equilibrium strategies for two negotiation procedures and the corresponding negotiation outcomes. For those outcomes, the authors further identify the conditions for which agents have identical preferences over the optimal agenda and procedure, and use those to reach agreements. The two negotiation procedures, issue-by-issue and package deal, are similar to ours. However, a key difference is that our initiating agent manages concurrent negotiations rather than *only one* bilateral negotiation. We also focus on possible issue switching rather than re-arranging the negotiation order. Our agents decide negotiation strategies based on the confidence in other agents' behavior in dynamic and uncertain environments, rather than the preferences or constraints as in [4]. Finally, in our framework, the agents learn over time the appropriate negotiation approach to improve negotiation outcome and reduce negotiation time cost in complex environments, instead of finding the optimal negotiation strategy (or outcome).

2.2. Combination of tactical negotiation and strategic negotiation

In complex negotiation environments where both negotiation tactics and negotiation management are critical, combining tactical negotiation and strategic negotiation is necessary. In such environments, tactical negotiations are needed for agents to conduct step-by-step negotiation tactics while strategic negotiations are needed to manage negotiations. Approaches are such as the preference-driven approach (e.g., [14]), the heuristic approach (e.g., [11]), the utility-based approach (e.g., [1]), and the syntax-dependent approach (e.g., [21]).We discuss these approaches in the following.

Sierra et al. [14] combine tactical negotiation and strategic negotiation using a preference-driven approach in order to conduct many-to-many (i.e., multiple initiating agents and multiple responding agents) negotiations on multiple issues. To achieve flexibility during negotiation, they set different weights for optional negotiation tactics, allowing the varying importance of the criteria to be modeled over time. For example, when the time resource is sufficient, the tactics that emphasize the behavior of other agents will be given greater preference than the tactics which base their values on the amount of time remaining. Thus

strategies combine tactics depending on the history of negotiations and the mental state of agents. Our work is similar to [14] in that both are concerned about time. However, we combine the strategy together with tactics based on the environmental change rather than the preferences, being more adaptive to the complex environments. The combination of strategic negotiation and tactical negotiation in our work is loosely coupled, which provides better modularity and flexibility.

Nguyen and Jennings [11] propose a heuristic approach to coordinate 1-to-many (i.e., one initiating agent and multiple responding agents) negotiations on a single issue. The initial strategy selection is based on the initiating agents' beliefs about others, and the negotiation tactics of agents in individual negotiations depend on the negotiation status of the same issue. Since they do not consider multi-issue negotiations, there is no issue assignment work and no issue will be switched from the pipeline of one agent to another one, which is a key consideration in our work.

Chevaleyre et al. [1] propose a utility-based approach to manage many-to-many negotiations on multiple issues. All agents involved in the concurrent negotiations model their preferences by means of tree-structured utility functions. An identical protocol is used to guide the multilateral negotiations by modifying the utility functions of each agent during negotiations. As a result of a sequence of structurally simple deals, an agreement with the maximal utilitarian social welfare (i.e., the maximal sum of all pleasures) can be achieved. In our work we address the 1-to-many negotiations rather than the many-to-many negotiations. The objective of our strategy is to improve negotiation outcome and lower negotiation cost, rather than simply maximize the negotiation outcome. In our complex environments, the means of tree-structured utility functions is not appropriate because it is time-consuming for agents to update the complex utility functions adaptive to the dynamic and uncertain change of the environments.

Zhang and Zhang [21] propose a negotiation solution based on the syntax-dependent formalism of belief revision to address the bilateral negotiation. They combine the conflict resolution strategy together with the step-by-step negotiation tactics. Specifically they divide a negotiation procedure into two stages. In the first stage, two agents meet together to work out mutually acceptable core demand from each side by step-by-step. In the second stage, each agent adjusts its own demand in order to make it consistent with the consent reached in the first stage by conducting a course of belief revision. The outcome of the negotiation is then the intersection of two agents' revised demand sets. The combination of the two types of negotiations in their work is loosely coupled. Different from their work, we address the 1-to-many negotiations rather than the 1-to-1 negotiations. We use strategic negotiation to conduct effective and efficient concurrent negotiations rather than simply resolve conflicts.

3. Framework

Our overall framework is based on a model called the Multi-Phase Coalition Formation (MPCF) model. The model consists of three phases: coalition planning, coalition instantiation and coalition evaluation, as depicted in Figure 1. In *coalition planning*, the agent applies case-based reasoning (CBR) to obtain a coalition formation plan. In *coalition instantiation*, the agent carries out the planned formation strategy through 1-to-many negotiations with coalition candidates. In *coalition evaluation*, the agent evaluates the coalition formation process, the formed coalition structure (if a coalition is successfully formed), and the coalition execution outcome (if the coalition is executed eventually) to determine the utility of the planned strategy and reinforces the strategy. In the following, we briefly outline the MPCF model. For details on this model, please refer to [9,15,16]. For details on the tactical negotiations, please refer to [10].

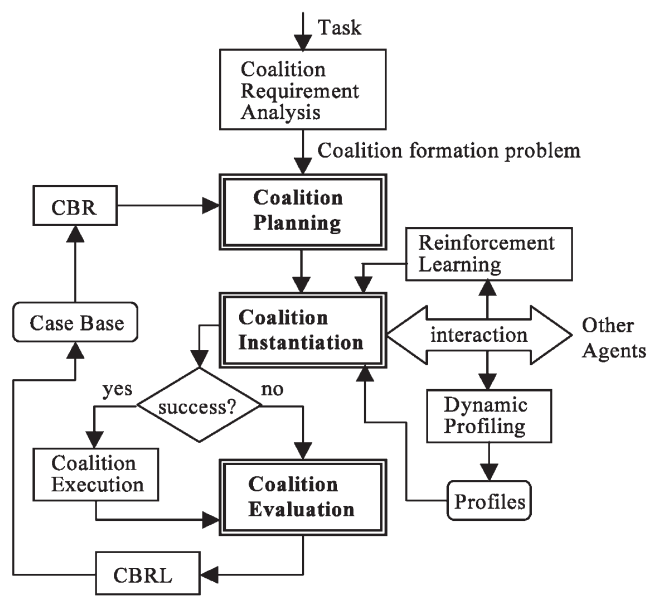

Figure 1. The Multi-Phase Coalition Formation (MPCF) model.

3.1. Coalition formation phases

In *coalition planning*, the coalition-initiating agent applies CBR to derive a specific coalition formation plan for the current problem based on a previous plan stored in the case-base. Given a problem (a task) to solve, the agent retrieves from its case-base the case with the most similar problem description to the current problem description. The agent then adapts the case *solution*, based on the differences between the two problem descriptions, to compose a coalition formation plan (see Section 3.2). Briefly, a plan is one that specifies guidelines for the agent on how to form the coalition, including the number of neighbors to approach, the number of coalition candidates to maintain, time allocated for the formation process, and so on.

The *coalition instantiation* phase implements the coalition formation plan to form a coalition. At first, the coalition-initiating agent normalizes the task— dividing the task into separate execution units as different negotiation issues, computing the potential utilities of its peers, and ranking the peers based on the potential utilities. Then the agent concurrently negotiates with each selected peer agent on the set of subtasks in an attempt to form the intended coalition. Each negotiation is argumentative where the initiating agent attempts to persuade the responding agent to perform a task or provide a resource by providing support or evidence for its request [18].

The *coalition evaluation* phase provides the basis for an agent to improve its coalition formation plans. This phase evaluates both the coalition instantiation process (in terms of time spent, number of messages used, number of peers approached, etc.) and the execution outcomes of the subtasks agreed upon in the coalition (in terms of the number of subtasks performed by highly-capable peers, etc.). In general, a good plan is one that uses little computational and communication resources with successful instantiations and subsequent executions.

3.2. Case-based reinforcement learning

We employ an integrated case-based reinforcement learning strategy [16] to utilize the results of the evaluation phase to improve the coalition planning phase. As alluded to in the previous section, our coalition formation framework relies on CBR to derive plans. Thus, the experience (or lessons learned) is stored in the cases.

A *case* consists of a problem description, a solution, an outcome, and its utility. The *problem* description consists of an agent's external and internal environments and the task description. The *solution* gives a coalition formation plan which specifies (1) the number of coalition candidates, (2) the number of expected coalition members, (3) the time allocated for coalition instantiation, (4) the task allocation, and (5) the number of messages recommended. The *outcome* documents the coalition instantiation results among agents, subtasks' execution results, and the evaluation values of the actual coalition formation process. The *utility* indicates the quality of the case, specifically, the quality of the plan in addressing the coalition problem represented in the case. Coupling the evaluation and the problem description, the agent can learn a new coalition formation case to increase its coverage of cases *or* can update the case's utility using the evaluation result to *reinforce* the case.

4. Confidence-based negotiation strategy

We first proposed the adaptive, confidence-based negotiation strategy in [17]. We employ this strategy to facilitate the initiating agents to conduct concurrent, strategic negotiations, during the coalition instantiation phase. Here we first present the motivations behind such a strategy in Section 4.1. Then, we describe the assumptions about the multiagent environment that are necessary for the strategy to be beneficial in Section 4.2. Finally, in Section 4.3, we discuss the algorithm for the strategy in details.

4.1. Strategy rationale

The objective of applying the proposed confidence-based negotiation strategy during the coalition instantiation phase is to help agents negotiate and form coalitions more effectively and more efficiently. It does not deal with how agents take step-by-step negotiation tactics; instead, it manages how the concurrent negotiations should be conducted in terms of scheduling and assigning different issues to different coalition candidates (here an issue refers to a subtask to be executed, especially the resource or capability required for the execution of that subtask).

Let us first describe the two strategic negotiation approaches: (1) *pipelined*, and (2) *packaged*. To help us illustrate the differences between the two approaches, suppose that at the end of the coalition planning

phase, the plan calls for the initiating agent a_i to approach two coalition candidates a_j and a_k (i.e., agents that have been deemed to likely be able to help) to negotiate over several subtasks for task T_m. Specifically, suppose a_j is allocated two subtasks T^1_m and T^3_m, and a_k is allocated three subtasks T^2_m, T^4_m, and T^5_m.

In the *pipelined approach*, a_i concurrently negotiates with a_j and a_k over *one issue at a time*, starting with T^1_m and T^2_m, respectively. If the negotiation between a_i and a_j completes successfully, then a_i initiates another negotiation with a_j to discuss T^3_m, and so forth. Further, suppose that a_i has successfully negotiated with a_j to handle T^1_m and T^3_m. At this point, if a_i realizes that (1) a_j is capable of performing T^5_m, and (2) the negotiation between a_i and a_k is still over T^4_m, then a_i can re-allocate T^5_m to a_j and start another negotiation. Similarly, if a negotiation over an issue fails, a_i will attempt to re-allocate that issue to another agent to try to salvage the coalition instantiation phase. That is, agent a_i has the flexibility of dynamically changing the task set selected for negotiations with a_k.

In the *packaged* approach, a_i packages multiple issues into one single negotiation. Using the same example above, that means a_i will initiate one negotiation with a_j, going over T^1_m and T^3_m in one shot, and another negotiation with a_k, trying to resolve all three issues, T^2_m, T^4_m, and T^5_m, at once. If two agents cannot reach an agreement on an issue, then the package to which the issue belongs fails as well. The last characteristic indicates that a package is either to be accepted (after negotiation) or to be rejected, and cannot be accepted piecewise issue by issue.

In complex environments that are dynamic, uncertain, and real-time, agent behaviors are difficult to perceive or predict accurately by other agents. For example, an initiating agent a_i may rank neighbor a_j as the best candidate and expect to reach an agreement for its negotiation in a short time. However, it is possible that a_j is busy and unable to entertain the negotiation request at that moment. Without a flexible management strategy, a_i would have to wait until a_j is available, probably missing the time deadline requirement. It is also possible that, because of the dynamic nature that we assume of our environment, the ranking of a neighbor by an agent may change during a negotiation and may thus require the agent to terminate the ongoing negotiation in favor of another. Thus we realize that there is a need for a management strategy that is flexible, capable of adapting to the profiled behavior of the neighbors as well as the real-time observation of the negotiation activities.

Comparing the two negotiation approaches, we see that the pipelined approach is more cautious, flexible, and opportunistic than the packaged approach. In the former, the initiating agent may dynamically redistribute an issue from the original pipeline to a "free and capable" pipeline of another coalition candidate. So the pipelined approach is preferred in uncertain and dynamic environments to increase the likelihood of successful negotiations. On the other hand, packaging multiple issues into one negotiation may reduce the number of agents to approach and subsequently decrease computational and communication overheads. So the packaged approach may be preferred in resource-constrained, real-time environments. An initiating agent thus needs to determine when to use which negotiation approach with which responding agents, for different issues during the same coalition formation task.

Note also that an initiating agent in our framework decides which approach to take based on the confidence it has in what it has learned of each of its neighbors. This allows the agent to consider volatility or uncertainty in the environment under resource constraints.

4.2. Assumptions on problem domain

Our adaptive, confidence-based negotiation strategy is based on the following three assumptions:

Assumption 1 (Overlapping Capabilities Assumption). We assume that in the multi-agent system, the agents have *overlapping* capabilities. This assumption facilitates the pipelined negotiation approach. If no agents share the same capabilities, then the pipelined approach will not have the flexibility to be opportunistic, or the possibility to exploit the responding agents that have completed negotiations early. For example, if there is only one candidate that knows how to perform a subtask T^1_m, then the agent a_i has no choice but to reserve T^1_m for that candidate.

Assumption 2 (Efficient Multi-Issue Negotiation Assumption). We assume that packaging multiple issues into a negotiation is more efficient than negotiating each issue one by one. That means the resource cost required by the packaged negotiation is less than that by pipelined negotiation. Different from the pipelined, one-at-a-time negotiation—in which an initiating agent needs to send a request and a responding agent needs to inform the negotiation outcome issue-by-issue, in packaged negotiations, the agents request or inform

over all the issues at one time. This assumption basically indicates that by packaging issues into one single negotiation an initiating agent is able to reduce the negotiation overhead of the same set of issues, which in turn reduces the computational and communication costs (e.g., the coalition formation time and the communication bandwidth).

Assumption 3 (Additive Multi-Issue Negotiation Arguments Assumption). We assume that to persuade a candidate (responding agent) to agree with a set of issues, the amount of arguments (or messages) that the initiating agent has to provide for all the issues in the packaged approach is the same as the total amount of all the arguments that it has to provide for each of the issues separately in the pipelined approach. This assumption supplements Assumption 2 in that the arguments needed by each issue in the packaged approach are the same as those in the pipelined approach. But the overhead of the former is less than the latter. In some scenarios, this assumption may not hold as negotiating multiple issues altogether may increase flexibility in reaching a compromise agreeable to both parties, and thus reducing the amount of arguments needed, at the expense of added reasoning capacity on the agents.

Assumption 4 (Inter-Coalition Competition Assumption). We assume that there is *inter-coalition* competition among coalition formation tasks. The inter-coalition competition may occur in different situations. First, *concurrent coalition formation tasks compete for the communication resources of the same agent*. Particularly, the same agent can be (i) a coalition-initiating agent that initiates multiple, concurrent coalition formation processes for these tasks, (ii) a responding agent that is approached for these concurrent tasks by different initiating agents, or (iii) a responding agent that is approached for these concurrent tasks of the same initiating agent. Second, *coalition formation tasks compete for the same capability (or problem-solving resource) of one same agent*. Particularly, the tasks can be requested by different initiating agents, or those by the same initiating agent, and overlap temporarily or functionally. Stemming from this intercoalition competition assumption, we assume that each task has a priority such that an agent will be able to choose to perform the task or subtask most important to itself.

Assumption 5 (Intra-Coalition Competition Assumption). We assume that there is *intra-coalition* competition among different subtask (negotiation issues) of the same coalition formation task. The intracoalition competition may occur in different situations: (1) when the coalition-initiating agent decides on the negotiation order of different subtasks that have been allocated to different neighbors, and (2) when the coalition-initiating agent requests for help from multiple agents for the same subtask. As a result of this assumption, we also assume that initiating agent will be able to sequence the negotiation issues according to their specified execution start times. We also further assume that the initiating agent has a degree of *greediness* that it exercises as an "insurance" policy. That is, if a subtask is of high priority, then to make sure that it will be able to get it solved, the initiating agent may choose to request from multiple agents that it knows to help solve the subtask. The "greedier" the initiating agent is, for example, the more agents it will approach.

4.3. Strategy description

The proposed adaptive, confidence-based negotiation strategy works as follows. First, an initiating agent a_i selects only agents that it knows to be capable of performing the subtasks of a task to be solved as coalition candidates. Second, depending on the subtasks and the capabilities of the coalition candidates, a_i allocates the subtasks as negotiation issues with specific coalition candidate. These first two steps are as dictated by the derived coalition formation plan as a result of casa-based reasoning. Third, during the coalition instantiation phase, a_i decides to use either the packaged or pipelined approach for each candidate based on its *confidence* in its perception of each coalition candidate's capabilities. For example, if the behavior of a coalition candidate a_j has been consistent as observed by a_i—e.g., a_j has always agreed to a_i's requests very quickly—then a_i can anticipate that negotiating with a_j over the new issues will likely lead to a high-utility outcome. However, if the behavior of a_j has not been consistent as observed by a_i, then a_i may choose to be cautious and do not package all issues into one negotiation with a_j.

An initiating agent a_i computes its confidence value in a candidate a_j in the following manner:

$$Conf_{a_i}(a_j) = \frac{\ell}{\hat{\ell}} \left[\begin{array}{l} \omega_{Chp} \sum_m \frac{\omega_m}{1+\sqrt{\sum_{s=1}^{\ell} (Chp^m_{a_i}(a_j,t_s) - \overline{Chp^m_{a_i}(a_j)})^2/(\ell-1)}} \\ + \omega_{Cho} \sum_n \frac{\omega_n}{1+\sqrt{\sum_{s=1}^{\ell} (Cho^n_{a_i}(a_j,t_s) - \overline{Cho^n_{a_i}(a_j)})^2/(\ell-1)}} \end{array} \right]$$

where $Chp^m_{a_i}(a_j, t_s)$ is the value of the a_j's mth characteristic related with coalition *process*, observed by a_i at time t_s ($s \in [1, \ell]$); and $Chp^n_{a_i}(a_j, t_s)$ is the value of one of a_j's nth characteristic related with coalition *outcome*, observed by a_i at time t_s ($s \in [1, \ell]$). Further, $\overline{Chp^m_{a_i}(a_j)}$ and $\overline{Chp^n_{a_i}(a_j)}$ are the corresponding means, and ℓ is the window size of a_i's observation periods ($\ell \geq 1$). A larger window size ℓ means more interactions between the two agents and thus may lead to higher confidence in a_i's observations of a_j. ω_m and ω_n are weights used to balance the values of the parameters related with coalition process and coalition outcome respectively. ω_{Chp} and ω_{Cho} are the weights used to emphasize the role of the process or the outcome. A higher ω_{Chp} is preferred in highly resource-constrained, real-time environments to facilitate a negotiation approach that may incur less resource costs (e.g., time); on the other hand, a higher ω_{Cho} is preferred in highly uncertain environments to have a higher likelihood of securing a successful negotiation outcome.

Note that a_i profiles a_j's behavior with two sets of characteristics, the process set (*Chp*) and the outcome set (*Cho*). The process set pertains to the negotiation process and consists of (1) the *tardiness* degree of the candidate a_j, indicating the time it takes a_j to respond to a request by a_i during negotiation (i.e., the average message round-trip-time between the two agents as observed by a_i), (2) the *hesitation* degree of the candidate a_j, indicating how ready a_j is to agree to a request by a_i (i.e., the amount of arguments that a_i needs to provide to persuade a_j), and (3) the *busyness* degree of the candidate a_j (i.e., the number of on-going negotiations between a_i and a_j at present). The outcome set pertains to the quality of the coalition outcome and consists of (1) the *satisfaction* degree of the candidate a_j, indicating how much a_j agrees to meet a_i's request (i.e., the negotiation utility of a_j to a_i for a particular issue (or subtask)), and (2) the *capability* of the candidate a_j, indicating how well a_j is capable of executing the agreed subtask (i.e., the execution utility of a_j to a_i on a particular issue (or subtask)). Note that the satisfaction degree is the upper bound on how well a subtask can be executed while capability is the actual observed quality of outcome. Due to the uncertainty in the environment, it is possible for the execution utility to be significantly less than the negotiation utility of a candidate to the initiating agent.

Given the above definition of confidence, an initiating agent a_i can choose to perform the pipelined approach with a candidate a_j if $Conf_{a_i}(a_j)$ less than a threshold κ, and to perform the packaged approach otherwise. A low κ may be preferred to avoid being overly conservative and always negotiating with each candidate in the pipelined approach if significant gains can be obtained from the packaged approach. For example, if the packaged approach significantly reduces the amount of arguments needed and therefore the time needed to negotiate successfully, then one may choose a low κ. On the other hand, in situations where task allocations are very dynamic and flexible, a high κ may be preferred to avoid being overly rigid and always negotiating with each candidate in the packaged approach.

Ultimately, the coalition formation process is based on a plan and the plan may not work as expected during the coalition instantiation process. Focusing only on the packaged approach vs. the pipelined approach, we see that it is possible for the initiating agent to lose contact with a candidate due to communication loss, agent faults, or extremely long communication delay. In such a case, the initiating agent needs to have a contingency plan to persist with the coalition instantiation process. In general, if the packaged approach is used, and a candidate is found to be non-responsive, then the agent may break up the package and distributes the issues to the remaining candidates in the coalition. For the pipelined approach, the response is more straightforward—simply transferring the issues to other pipelines. And, in the case where the initiating agent is stuck with a particular issue negotiating with a non-responsive candidate, the agent will duplicate another negotiation with another candidate on the same issue and conduct the two negotiations in parallel. Whichever negotiation completes successfully first will prompt the agent to terminate the other.

5. Experimental results

We implemented the adaptive, confidence-based negotiation strategy into the Multi-Phase Coalition Formation (MPCF) framework in a multiagent system. In the system, each agent has multiple overlapping capabilities and is capable of performing multiple tasks. When an agent encounters a task, it first analyzes whether it is able to solve the problem all by itself; if not, it initiates a coalition formation process. Each agent has 3 + N threads: (1) a *core* thread to manage tasks, reason, and learn, (2) a *communication* thread, (3) an *execution* thread for task simulation, and (4) N *ne-*

Table 1. The experiment sets

Experiment Sets	Dynamic Range	Uncertainty Range	Real-Time Degrees
EB1	$[RTT_{min}, 2\ RTT_{min}]$	[0.6, 1.0]	0 ~ 6
EB2	$[RTT_{min}, 2.5\ RTT_{min}]$	[0.4, 1.0]	0 ~ 6
EB3	$[RTT_{min}, 3\ RTT_{min}]$	[0.2, 1.0]	0 ~ 6

gotiation threads for concurrent negotiations with other agents.

In this paper, we report some experimental results aimed to quantitatively evaluate the adaptive, confidence-based negotiation strategy in complex environments by comparing its performance with a purely packaged and a purely pipelined strategy. The purely packaged strategy is one where the initiating agent packages all issues allocated to each candidate and negotiates over them via one negotiation process. If the two agents cannot agree on any of the issues, the entire package is considered a failure. The initiating agent then approaches another candidate that is free and capable of addressing all the issues of the package. The purely pipelined approach is one where all issues are individually negotiated, and issues can be re-distributed to available pipelines as long as the candidates are capable of addressing those issues.

We evaluate the performances in terms of coalition formation efficiency and effectiveness. The efficiency is measured in terms of the cost of the coalition formation process, while effectiveness is measured in terms of the quality of the outcome of the formed coalition.

5.1. Experiment sets

We have three experiment sets: EB1, EB2, and EB3, as shown in Table 1. There are three environment parameters that we vary for our experiments:

(1) *Dynamic range*: This represents the range within which agents' dynamic behavior varies. We have simulated a variety of dynamic ranges of the environment through setting different variation ranges of agents' tardiness degrees, as defined previously in Section 4.3. Thus, this range describes a key characteristic of the negotiation *process*. As shown in Table 1, EB3 is the most dynamic while EB1 is the least dynamic.

(2) *Uncertainty range*: This represents the range within which negotiation outcomes may vary given the same pair of agents and the same issues. Specifically, we simulate this range by varying agent's satisfaction degrees, as defined in Section 4.3 above. This characteristic has a direct impact on negotiation *outcome*. As shown in Table 1, a 1.0 means that the satisfaction degree is 100% (a full deal can be reached), a 0.2 means that it is only 20% (only a partial deal can be reached). Thus, EB3 is the most uncertain while EB1 is the least uncertain.

Table 2. Real-time environmental settings

Real-Time Degrees	Time Limits (ticks)
0	50000
1	27500
2	25000
3	22500
4	20000
5	17500
6	15000

(3) *Real-time degree*: This represents the time constraints in negotiations. Any negotiation that runs longer than its time limit is aborted, yielding a 0 for the outcome utility. Table 2 shows the actual time limit for each real-time degree. A degree of 0 indicates that the environment is non-real-time as the agents will have more than sufficient time to carry out their negotiations. Thus, a degree of 1 is the least time-constraining while a degree of 6 is the most time-constraining. All three experiment sets have the same range of time constraints.

Note that due to the real-time constraints, it is possible that an agent never gets around to tackle all subtasks of a task. Thus, there will be subtasks that are not yet negotiated at the end of the time limit. For these subtasks, each will count as 0 towards the overall utility of the task. In this manner, we not only keep track of failed negotiations, but also take into account negotiations waiting in the task queue that an agent fails to initiate.

5.2. Efficiency: Cost in negotiation processes

To examine whether the confidence-based strategy outperforms others in negotiation cost, we compare each strategy's average number of messages sent for negotiations per task, and average negotiation time per task in non-real-time environments (ES1, ES2, and ES3). We specify the *negotiation time* of a coalition

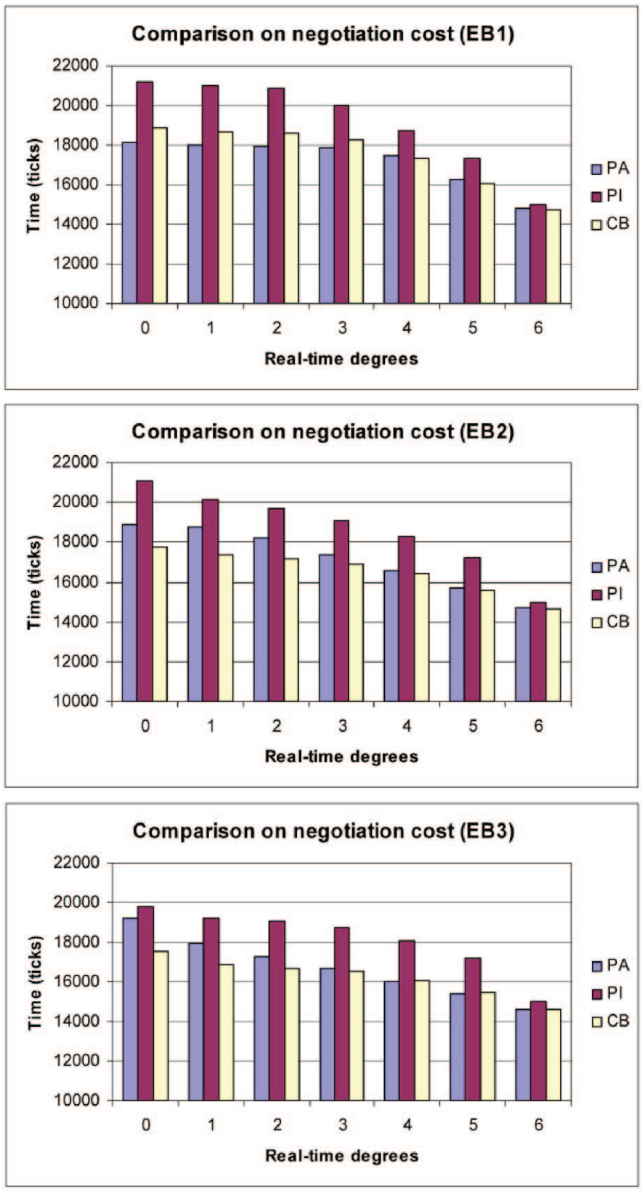

Figure 2. Comparison on negotiation time in EB1, EB2, and EB3, from the less dynamic and uncertain to the more dynamic and uncertain environments. PA = purely-packaged, PI = purely-pipelined, and CS = confidence-based.

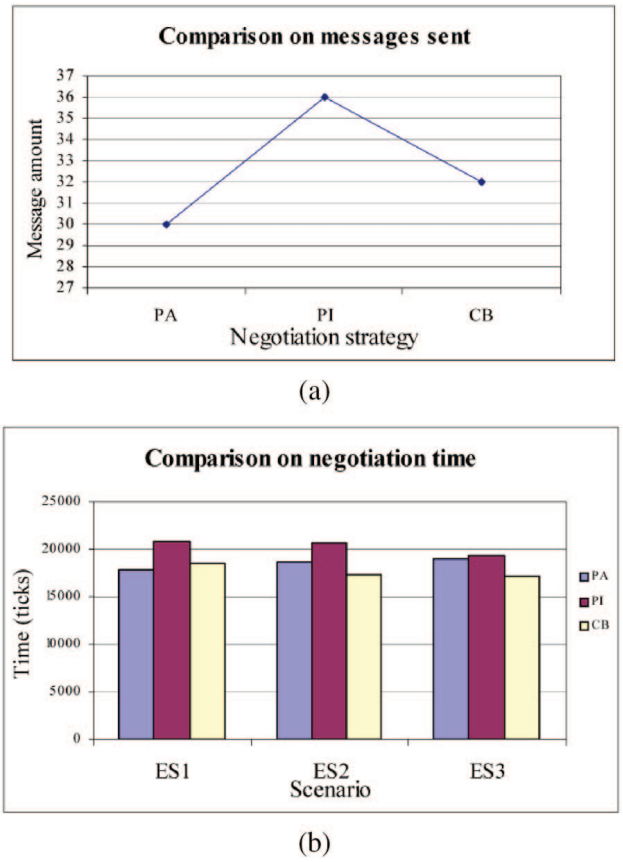

Figure 3. Comparison on negotiation cost: (a) message cost, and (b) time cost. PA = purely-packaged, PI = purely-pipelined, and CS = confidence-based.

formation task from the moment when a coalition-initiating agent begins to negotiate for the task to the moment when the agent receives the last outcome of negotiations. If the time limit comes first before all issues (subtasks) are negotiated, then all ongoing negotiations will be aborted, and remaining issues that are yet-to-be-negotiated will also be removed. In this case, the negotiation time of the task is the time limit allocated to the task. Figure 2 shows the results over different real-time degrees, where PA represents the purely-packaged negotiation strategy, PI represents the purely-pipelined negotiation strategy, and CB represents the adaptive, confidence-based negotiation strategy, while Figure 3 shows the averaged results of the three strategies.

In non-real-time environments, the number of messages sent for negotiations in each task depends primarily on (1) the number of issues in the task, (2) the number of arguments required by each candidate, (3) the number of candidates approached with the packaged or the pipelined strategy, and (4) the issues' distribution pattern among the candidates. The environment's dynamism or uncertainty is relatively a nonimpact, since the agents have ample time to carry out their negotiations and to re-distribute issues to obtain the best deals. From Figure 3(a), we observe that the confidence-based strategy needed more messages than the purely packaged, but fewer messages than the purely pipelined. This result is expected. First, with Assumption 3, the only way for the packaged approach to outperform the pipelined approach is its ability of lowering message overhead as agents using the packaged approach deal with fewer requests. As a result, PA outperforms CB, and CB outperforms PI.

From Figs 2 and 3(b), we observe that in all three experiment sets, the average negotiation time of tasks in PI is always the highest. Note that in CB, the agents choose when to use a mixture of packaged and pipelined approaches when forming coalitions. Since it is evident that the CB outperforms PI, that implies that the agents are able to make the correct decisions—picking the right candidates to use the pipelined or packaged approach—to improve their efficiency.

Furthermore, in non-real-time environments (i.e., when real-time degree is 0), when the environment is more dynamic and uncertain (moving from EB1 to EB3), the average negotiation time for PA increases, while it decreases for both PI and CB. Upon closer analysis, we realize that when the environment is more dynamic and uncertain (as in EB2 and EB3), it is more likely for a candidate to possibly respond more slowly. That results in more chances for some issues waiting in the pipelines to be moved from slow-progressing pipelines to the fast-progressing ones, if the latter have empty slots in the schedule of the pipelines. The non-critical time constraint also makes these switches possible. This thus results in the lower values of negotiation times observed for PI and CB. This implies that in more dynamic and uncertain environments, a design that incorporates some pipelined approach can outperform a purely packaged one in negotiation efficiency.

Interestingly, in EB1, EB2, and EB3, when the real-time degrees of the environments are increased, the difference in the average negotiation time between CB and PA decreases. Especially when the real-time degree increases from 4 to 6, the average negotiation time in the confidence-based strategy is almost equal to that in the purely-packaged strategy. Upon further analysis, we find that this is because when the time constraints became more critical, more negotiations were aborted sooner so much so that moving issues around from one candidate to another was futile: there are simply not enough agents to handle all the necessary subtasks or issues within the time limit. As a result, the difference is insignificant in negotiation time between CB and PA.

Therefore, we see that the adaptive, confidence-based negotiation strategy can outperform the purely packaged strategy in negotiation cost when the environment is less real-time, and more dynamic and uncertain; but cannot when the environment is less dynamic and uncertain. On the other hand, the confidence-based strategy generally outperforms the purely pipelined strategy in terms of the average negotiation time in all non-real-time scenarios. That the agents using the confidence-based strategy are able to save time using packaged negotiations indicates that the environment is still consistent enough for the agents to make confident and useful profiles of their interactions with other agents. This also indicates that our experiments do not identify a cutoff point above which the purely pipelined strategy will always outperform the confidence-based strategy. Ultimately, we hypothesize that when the environment is extremely dynamic and uncertain, the confidence-based strategy will conform to the purely pipelined one.

5.3. Effectiveness: Negotiation outcome

To examine whether the confidence-based strategy outperforms others in the quality of the negotiation outcome, we compare each strategy's average negotiation outcome utility of all the tasks in the experiment sets, shown in Figure 4. Note that the outcome utility reported is the average of all 99 tasks negotiated, and each task's negotiation outcome utility is the average of the satisfaction degree achieved for each of its issues (subtasks).

From Figure 4, we observe that in non-real-time environments (real-time degree = 0), in EB1, EB2, and EB3, the difference in the negotiation outcome among the three approaches is insignificant. That is, with sufficient negotiation time, the more flexible PI approach does not yield better negotiation outcome than does the confidence-based strategy or the purely-pipelined strategy. Furthermore, when the environment is less time-critical, especially when the environment is less dynamic and uncertain (e.g., in EB1, real-time degree is from 1 to 3), CB still fails to outperform PA. On one hand, the coalition-initiating agent in CB is more confident in its coalition candidates that it selects the packaged approach more often. On the other hand, with relatively sufficient negotiation time and relatively fast response time, the pipelined approach does not yield better average negotiation outcome for the confidence-based strategy. However, CB does perform as well or outperforms both PA and PI in all but one (i.e., EB1, real-time degree = 3) real-time environments, validating the confidence-based strategy that adaptively combines the more opportunistic pipelined approach and the more efficient packaged approach to conduct multiple, concurrent negotiations effectively.

Note that as the real-time constraints become more stringent, the negotiation outcomes in all the strategies deteriorate, as expected since there are more negotiations

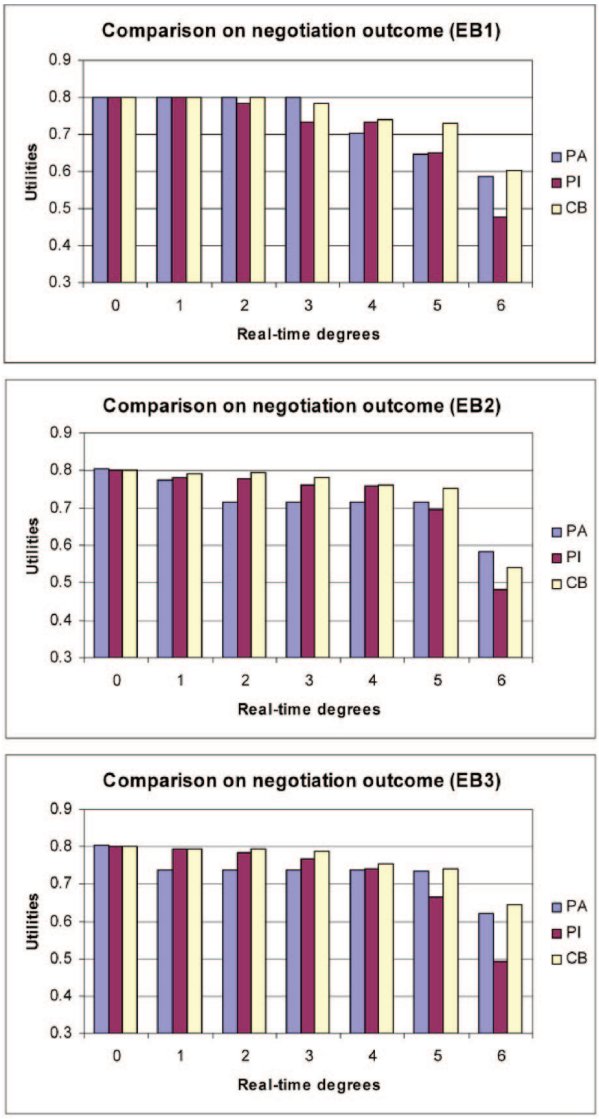

Figure 4. Comparison on negotiation outcome in EB1, EB2, and EB3, from the less dynamic and uncertain to the more dynamic and uncertain environments. PA = purely packaged, PI = purely pipelined, and CB = confidence-based.

aborted due to the time limits. Noticeably, CB's performance deteriorates at the slowest rate. This hints that (a) the confidence-based strategy is more robust than PA and PI, and (b) the strategy is able to make good decisions on when to use the packaged or pipelined negotiation approach appropriately.

Unexpectedly, PA outperforms PI in terms of average negotiation outcome in most scenarios, especially when the real-time constraint is high. This is counter-intuitive because the pipelined negotiation approach is supposed to be more flexible and opportunistic. Upon closer analysis, we realize that when a negotiation task is highly time-constrained, PI may not be able to negotiate all issues. In that case, being opportunistic does not necessarily pay off: if the negotiations are not completed soon enough, the issues waiting in the pipelines will not get a chance to be moved from a slower pipeline to a faster one. Thus, the advantage of the pipelined approach lessens. On the other hand, by packaging issues into one single negotiation, the purely-packaged strategy can actually by chance complete negotiations within the imposed time limit for some tasks. Thus, we have found a cutoff point in terms of real-time degree above which the purely-packaged strategy starts to outperform the purely-pipelined strategy for the three different scenarios.

Based on the above results and analysis, we see that (1) the adaptive, confidence-based negotiation strategy can outperform the purely-pipelined strategy in negotiation outcome when the environment is time constrained; and (2) the adaptive, confidence-based negotiation strategy can outperform the purely-packaged strategy in negotiation outcome when the environment is highly time-constrained but less dynamic and uncertain, or when the time-constrained environment is more dynamic and uncertain.

6. Conclusions and future work

We have described an adaptive, confidence-based negotiation strategy in the problem domain of multiagent coalition formation. This strategy addresses multi-issue strategic negotiations in complex environments. In our strategy, an agent chooses a combination of pipelined and packaged approaches to conduct multiple, concurrent negotiations with its peers based on the agent's confidence in its peer's consistency. Our experiments have found that the confidence-based strategy generally outperforms the purely packaged and purely pipelined strategies in a variety of dynamic, uncertain, and real-time environments. We have also found a cutoff point where the strategy starts to lose ground to the purely packaged approach. Further, we have observed how the purely packaged approach can be more effective than the purely pipelined one in a highly dynamic and uncertain environment, when it is highly time-constrained. Our results are also supported by [2]. Note that the computation of the confidence measure can be seen as a form of learning: an initiating agent's assessment of a candidate's capability gained from a series of observations over time. Excelente-Toledo and Jennings [2] showed that learning improves the decision making when agents are uncer-

tain about the other agents' actions, and that learning is ineffective when agents operate in more static environments when compared to other agents that can make reasonably accurate predictions about their environment and other agents. A key difference is that our agents make decisions to use the packaged or pipeline approach based on their *confidence* in what they have learned, instead of making decisions based directly on what they have learned.

Our current and immediate work include conducting further experiments (1) with larger dynamism and uncertainty ranges to locate the cutoff point beyond which the purely pipelined strategy outperforms the confidence-based strategy, (2) with more agents and more issues (and more diverse capabilities) to study the rate of an agent in finally feeling the impact of the confidence-based strategy—presently, the experiments were setup so that an agent was able to profile a peer's consistency quite accurately after only a few tasks, and (3) with noise injected into the environment such as noise in communication (e.g., messages getting dropped) and the specification of tasks (e.g., issues change after receipt).

Further, we look to improving the negotiation strategy to address the uncertainty of the environment. How should the agents learn how to pick the right pipeline to which to move negotiation issues? Presently, the initiating agent simply opportunistically chooses the pipeline that is free and where the candidate is capable. However, a free pipeline and a capable candidate do not guarantee that the candidate will agree to handle the new issue. This is because, in a complex environment, an agent that is not busy now may become very busy in the next few moments, or a communication link that is fast now may become very congested and slow next. Based on the confidence that the initiating agent has in its candidates and past "transfer" utility, perhaps the agent can learn to be more conservative and patient against the uncertainty of the environment. This is the key to improve the quality of the coalition since the original assignment of subtasks (issues) to each candidate is deemed to be the best plan that the agent has derived during its coalition planning phase.

Acknowledgments—This work was partially supported by the Department of Computer Science and Engineering at the University of Nebraska, and by a subcontract through the University of Kansas on a DARPA-funded ANTS project. The authors would like to thank Nobel Khandaker for his review of the paper.

References

[1] Y. Chevaleyre, U. Endriss and N. Maudet, Tractable negotiation in tree-structured domains, in: *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'06)*, Hakodate, Japan, 2006, pp. 362– 369.

[2] C. B. Excelente-Toledo and N. R. Jennings, Learning when and how to coordinate, *Web Intelligence and Agent Systems* **1**(3) (2003), 203–218.

[3] S.S. Fatima, M. Wooldridge and N. R. Jennings, Multi-issue negotiation under time constraints, in: *Proceedings of the 1st International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'02)*, Bologna, Italy, 2002, pp. 143– 150.

[4] S.S. Fatima, M. Wooldridge and N. R. Jennings, Optimal negotiation of multiple issues in incomplete information settings, in: *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'04)*, New York NY, 2004, pp. 1080–1087.

[5] S. Kraus, *Strategic Negotiation in Multi-Agent Environments*, Cambridge MA, The MIT Press, 2001.

[6] S. Kraus and O. Schechter, Strategic-negotiation for sharing a resource between two agents, *Computational Intelligence* **19**(1) (2003), 9–41.

[7] S. Kraus and J. Wilkenfeld, A strategic negotiations model with applications to an international crisis, *IEEE Transactions on Systems, Man, and Cybernetics* **23**(1) (1993), 313–323.

[8] S. Kraus, K. Sycara and A. Evenchik, Reaching agreements through argumentation: A logical model and implementation, *Artificial Intelligence* **104**(1-2) (1998) 1–69.

[9] X. Li, *Improving Multiagent Coalition Formation in Complex Environments*, PhD Dissertation, Computer Science and Engineering Department, University of Nebraska, Lincoln, NE, 2007.

[10] X. Li and L.-K. Soh, Hybrid Negotiation for Resource Coordination in Multiagent Systems, in: *Web Intelligence and Agent Systems (WIAS)* **3**(4) (2005), 1–30.

[11] T. D. Nguyen and N. R. Jennings, Coordinating multiple concurrent negotiations, in: *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'04)*, New York NY, 2004, pp. 1064–1071.

[12] H. Raiffa, *The Art and Science of Negotiation*, Cambridge MA, Harvard University Press, 1982.

[13] J. Rosenschein and G. Zlotkin, *Rules of Encounter Designing Contentions for Automated Negotiation among Computers*, Cambridge MA, The MIT Press, 1994.

[14] C. Sierra, P. Faratin and N. R. Jennings, A service-oriented negotiation model between autonomous agents, in: *Proceedings of the 8th European Workshop on Modeling Autonomous Agents in a Multi-Agent World (MAAMAW'97)*, Ronneby, Sweden, 1997, pp. 17–35.

[15] L.-K. Soh and X. Li, A learning-based coalition formation model for multiagent systems, in: *Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'03)*, Melbourne, Australia, 2003, pp. 1120–1121.

[16] L.-K. Soh and X. Li, An integrated multilevel learning approach to multi-agent coalition formation, in: *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI-03)*, Acapulco, Mexico, 2003, pp. 619–624.

[17] L.-K. Soh and X. Li, Adaptive, confidence-based multiagent negotiation strategy, in: *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS'04)*, New York NY, 2004, pp. 1048–1055.

[18] L.-K. Soh and C. Tsatsoulis, Reflective negotiating agents for real-time multisensor target tracking, in: *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI-01)*, Seattle WA, 2001, pp. 1121–1127.

[19] K. Sycara, Argumentation: Planning other agents' plans, in: *Proceedings of the 11th International Joint Conference on Artificial Intelligence (IJCAI-89)*, Detroit MI, 1989, pp. 517– 523.

[20] D. N. Walton and E. C. W. Krabbe, *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*, Albany NY, SUNY Press, 1995.

[21] D. Zhang and Y. Zhang, A computational model of logic-based negotiation, in: *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI-06)*, Boston MA, 2006, pp. 728–733.